

Strategies for Global Optimization of Temporal Preferences

Content Areas: Temporal Reasoning, Constraint Satisfaction

Abstract

A temporal reasoning problem can often be naturally characterized as a collection of constraints with associated local preferences for times that make up the admissible values for those constraints. Globally preferred solutions to such problems emerge as a result of well-defined operations that compose and order temporal assignments. The overall objective of this work is a characterization of different notions of global preference, and to identify tractable sub-classes of temporal reasoning problems incorporating these notions. This paper extends previous results by refining the class of useful notions of global temporal preference that are associated with problems that admit of tractable solution techniques. This paper also answers the hitherto open question of whether problems that seek solutions that are globally preferred from a *utilitarian* criterion for global preference can be found tractably.

Introduction

Many temporal reasoning problems can be naturally characterized as collections of constraints with associated local preferences for times that make up the admissible values for those constraints. For example, one class of vehicle routing problems (Toth and Vigo 2001) consists of constraints on requested service pick-up or delivery that allow flexibility in temporal assignments around a specified fixed time; solutions with assignments that deviate from this time are considered feasible, but may incur a penalty. Similarly, dynamic scheduling problems (El Sakkout *et al.* 1998), whose constraints may change over time, thus potentially requiring solution revision, often induce preferences for revised solutions that deviate minimally from the original schedule.

To effectively solve such problems, it is necessary to be able to order the space of assignments to times based on some notion of global preference, and to have a mechanism to guide the search for solutions that are globally preferred. Such a framework arises as a simple generalization of the Simple Temporal Problem (STP) (Dechter *et al.* 1991), in which temporal constraints are associated with a local preference function that maps admissible times into values; the

result is called *Simple Temporal Problem with Preferences* (STPP) (Khatib *et al.* 2001). Globally optimal solutions to STPPs emerge as a result of well-defined operations that compose and order partial solutions.

Different concepts of composition and comparison result in different characterizations of global optimality. Past work has introduced three notions of global preference: Weakest Link (maximize the least preferred time), Pareto, and Utilitarian. Much of the work to date has been motivated by the overall goal of finding tractable solutions to temporal optimization problems with realistic global preference criteria. In particular, NASA is motivated to create systems that will automatically find optimally preferred solutions to problems in the rover planning domain (Bresina *et al.* 1997), where the goal is to devise plans for visiting a number of scientifically promising science targets.

In addition to reviewing the STPP framework (section 2), this paper extends previous results motivated by the overall goal of identifying useful notions of global preference that correspond to problems that can be solved tractably. First, we introduce a new category of global optimality called *stratified egalitarian* optimality, and prove that it precisely characterizes the subset of Pareto optimal solutions returned by a tractable technique introduced previously (section 3). Second, we provide an affirmative answer to the question of whether utilitarian optimal solutions to temporal reasoning problems can be also found tractably within this framework (section 4). This paper closes with a discussion of future work.

Simple Temporal Problems with Preferences

A *temporal constraint* depicts restrictions on the distance between arbitrary pairs of distinct events. In (Khatib *et al.* 2001), a *soft temporal constraint* between events i and j is defined as a pair $\langle I, f_{ij} \rangle$, where I is a set of intervals $\{[a, b], a \leq b\}$ and f_{ij} is a *local preference function* from $\bigcup I$ to a set A of admissible preference values. For the purposes of this paper, we assume the values in A are totally ordered, and that A contains designated values for minimum and maximum preference.

When I is a single interval, a set of soft constraints defines a *Simple Temporal Problem with Preferences* (STPP), a generalization of Simple Temporal Problems [Dechter, *et al.*, 1991]. An STPP can be depicted as a pair (V, C') where

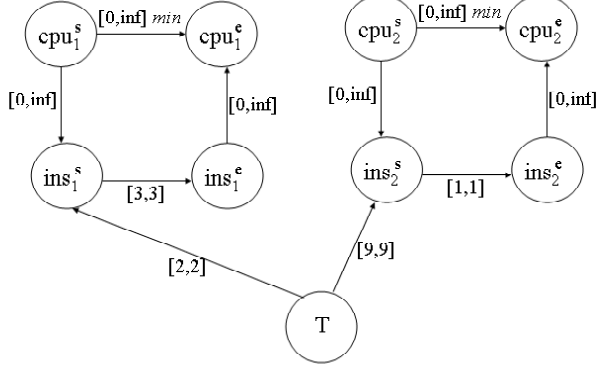


Figure 1: The STPP for a Rover Science Planning Problem (T is any timepoint)

V is a set of variables standing for temporal distances, and $C = \{[a_{ij}, b_{ij}], f_{ij}\}$ is a set of soft constraints defined over V . An STPP, like an STP, can be organized as a network of variables representing events, and links labeled with constraint information. A *solution* to an STPP is a complete assignment to all the variables that satisfies the temporal constraints.

We define a *preference vector* of all the local preference values associated with a set $F = \{f_{ij}\}$ of local preference functions and a solution S . Formally, let $f_{ij}(S)$ refer to the preference value assigned by f_{ij} to the temporal value that S assigns to the distance between events i and j , and $U_{F(S)} = \langle f_{12}(S), f_{13}(S) \dots f_{n-1n}(S) \rangle$ be the *preference vector* associated with F and S . In what follows the context will permit writing U_S instead of $U_{F(S)}$ without ambiguity, and U_S^k will refer to the k^{th} preference value of U_S .

For an example of an STPP, consider a simple Mars rover planning problem, illustrated in Figure 1. The rover has a sensing instrument and a CPU. There are two sensing events, of durations 3 time units and 1 time unit (indicated in the figure by the pairs of nodes labeled ins_1^s, ins_1^e and ins_2^s, ins_2^e respectively). The event T depicts a reference time point (sometimes referred to as “the beginning of time”) that allows for constraints to be specified on the start times for events. There is a hard temporal constraint that the CPU be on while the instrument is on, as well as a soft constraint that the CPU should be on as little as possible, to conserve power. This constraint is expressed in the STPP as a function from temporal values indicating the duration that the CPU is on, to preference values. For simplicity, we assume that the preference function min on the CPU duration constraints is the negated identity function; i.e., $min_{ij}(t) = -t$; thus higher preference values, i.e. shorter durations, are preferred.

A solution to an STPP has a *global preference value*, obtained by combining the local preference values using operations for composition and comparison. Optimal solutions to a STPP are those solutions which have the best preference value in terms of the ordering induced by the selected

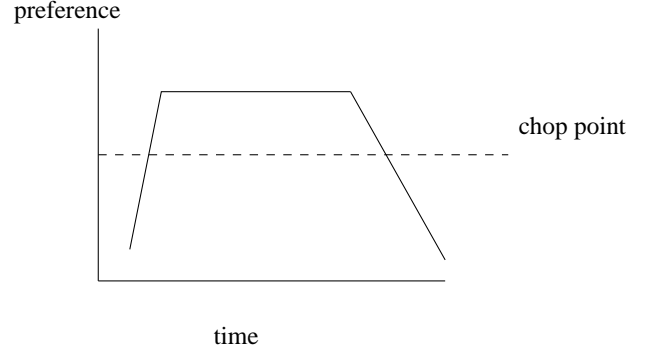


Figure 2: “Chopping” a semi-convex function.

comparison operator. Solving STPPs for globally preferred assignments has been shown to be tractable, under certain assumptions about the “shape” of the local preference functions and about the operations used to compose and compare solutions.

For example, first consider a class of local preference functions that includes any function such that if one draws a horizontal line anywhere in the Cartesian plane of the graph of the function, the set of X such that $f(X)$ is not below the line forms an interval. This class of *semi-convex* functions includes linear, convex, and also some step functions.

Second, consider an STPP solver based on the notion of *Weakest Link Optimization* (WLO). This framework consists of an operator for composing preference values in A based on the minimal value of the component values, and an operator for comparing preference values in A which returns the maximum of the two values. This framework induces an evaluation of solutions based on a single, “weakest link” value. Given preference vectors U_S and $U_{S'}$ corresponding to distinct solutions S and S' , we will say that S is *Weakest-Link-Optimal* (WLO) -preferred to S' , or S' is *WLO-dominated* by S , if $\min(U_{S'}) < \min(U_S)$, where $\min(U)$ returns the minimum value of vector U . *WLO-optimal solutions* are those to which no other solutions are WLO-preferred.

STPPs with semi-convex preference functions for WLO-optimal solutions can be solved tractably by a process called the *chop method*. This method is based on the act of “chopping” a preference function (Figure 2). Semi-convexity implies that the set of times for which the preference function returns a value above a selected chop point forms a convex interval; call this interval the *chop-induced constraint*. For a set of preference functions in an STPP, chopping all of them at the same preference value induces a Simple Temporal Problem, namely, of finding a set of assignments that satisfies all the chop-induced constraints. A binary search will return the largest preference value v_{opt} for which a solution to the induced STP exists; it can be shown that the solutions at v_{opt} are *WLO-optimal*.

Because the chop method returns the solution to an STP, its output is a *flexible temporal plan*, i.e., a set of solutions that have the same WLO-optimal value. Flexibility is often considered important in ensuring robustness in an execu-

tion environment that is uncertain (Muscettola *et al.* 1998). Nonetheless, the WLO criterion for globally preferred solutions has the disadvantage of being “myopic”, in the sense that it bases its evaluation on a single value. This feature can be shown to limit its usefulness in solving real temporal planning problems. The rover example in Figure 1 can be used to illustrate this myopia. Because the CPU must be on at least as long as the sensing events, any globally preferred solution using WLO has preference value -3. The set of solutions that have the WLO-optimal value includes solutions in which the CPU duration for the second sensing event varies from 1 to 3 time units (again, since WLO bases its evaluation solely on the least preferred value). The fact that WLO is unable to discriminate between the global values of these solutions, despite the fact that the one with 1 time unit is obviously preferable to the others, can be clearly viewed as a limitation.

Less myopic global preference criteria can be defined. For example, we can say that S' *Pareto-dominates* S if for each j , $U_S^j \leq U_{S'}^j$, and for some k , $U_S^k < U_{S'}^k$. The *Pareto optimal set* of solutions is the set of non-Pareto-dominated solutions. Similarly, we can say that S' *Utilitarian-dominates* S if $\sum_j U_S^j \leq \sum_j U_{S'}^j$, and the *Utilitarian optimal set* of solutions is the set of non-Utilitarian-dominated solutions.

In a previous result (Khatib *et al.* 2003), it was shown that a restricted form of Pareto-optimality can be achieved by an iterative application of the chop method. The intuition is that if a constraint solver could “ignore” the weakest link values (i.e. the values that determined the global solution evaluation) then it could eventually recognize solutions that dominate others in the Pareto sense. The links to be ignored are called *weakest link constraints*: formally, they comprise any link in which the optimal value for the preference function associated with the constraint is the same as the WLO value for the global solution. Formalizing the process of “ignoring” weakest link values is a two-step process of committing the flexible solution to consist of the interval of optimal temporal values, and reinforcing this commitment by resetting their preferences to a single, “best” value. Formally, the process consists of:

- Squeezing the temporal domain to include all and only those values which are WLO-optimally preferred; and
- Replacing the preference function to one that assigns the highest (most preferred) value to each element in the new domain.

The first step ensures that only the best temporal values are part of any solution, and the second step allows WLO to be re-applied to eliminate Pareto-dominated solutions from the remaining solution space. The resulting algorithm, called WLO+ returns, in polynomial time, a Simple Temporal Problem (STP) whose solutions are a nonempty subset of the WLO-optimal, Pareto-optimal solutions to a STPP. The algorithm WLO+ from (Khatib *et al.* 2003) is reproduced in Figure 3 for completeness. Where C is a set of soft constraints, the STPP (V, C_P) is solved (step 3) using the chop approach. In step 5, we depict the soft constraint that results from the two-step process described above as

Inputs: an STPP $P = (V, C)$

Output:

A STP (V, C_P) whose solutions are Pareto optimal for P .

- (1) $C_P = C$
- (2) while there are weakest link soft constraints in C_P do
- (3) Solve (V, C_P)
- (4) Delete all weakest link soft constraints from C_P
- (5) For each deleted constraint $\langle [a, b], f \rangle$,
- (6) add $\langle [a_{opt}, b_{opt}], f_{best} \rangle$ to C_P
- (7) Return (V, C_P)

Figure 3: STPP solver WLO+ returns a solution in the Pareto optimal set of solutions

$\langle [a_{opt}, b_{opt}], f_{best} \rangle$, where $[a_{opt}, b_{opt}]$ is the interval of temporal values that are optimally preferred, and f_{best} is the preference function that returns the most preferred preference value for any input value. Notice that the run time of WLO+ is $O(|C|)$ times the time it takes to execute $Solve(V, C_P)$, which is a polynomial.

WLO+, applied to the rover example in Figure 1, finds a Pareto optimal solution in two iterations of the while loop. In the first iteration, the weakest link is that between the start and end of the first CPU event. WLO+ deletes this link and replaces it with one with the interval $[3, 3]$ and the local preference function f_{best} . This new STPP is then solved on the second iteration, whereby the WLO-optimal solution with the CPU duration of 1 is generated. The solution to this STPP is a Pareto-optimal solution to the original problem.

WLO+ was a positive result in the search for tractable methods for finding globally preferred solutions based on less myopic criteria for global preference than WLO-optimality. We now proceed to refine and expand these results in two ways: first by offering a more concise characterization of the class of solution returned by WLO+, and secondly, by showing how restricted classes of STPP with a Utilitarian criterion for global preference can be solved tractably.

WLO+ and Stratified Egalitarianism

As noted in the previous section, the minimal network returned by running WLO+ on a STPP is a subset of the set of Pareto Optimal Solutions for that problem. In this section, we present a concise description of this set. By doing so, it is revealed that WLO+ is based on a useful concept of global preference.

We introduce a concept of global preference called Stratified Egalitarianism (SE). Consider again two preference vectors U_S and $U_{S'}$ associated with solutions S and S' . We will say S' SE-dominates S at preference level x if:

- $U_S^i < x$ implies $U_{S'}^i \geq U_S^i$.
- There exists an i such that $U_S^i < x$ and $U_{S'}^i > U_S^i$.
- $U_S^i \geq x$ implies $U_{S'}^i \geq x$.

We say that S' SE-dominates S (without further qualification) if there is any level x such that S' dominates S at x .

It is not hard to see that the SE-dominance relation is anti-symmetric and transitive¹, thus inducing a partial ordering of solutions. A solution S' will be said to be SE-optimal if it is not SE-dominated. Note that if a solution S' Pareto-dominates S , then S' SE-dominates S at the “highest” level of the S' vector. Thus, SE-optimality implies Pareto optimality. Furthermore, if S' is superior to S in the WLO ordering, then S' SE-dominates S at the “lowest” level of the S' vector. Thus, SE-optimality also implies WLO optimality.

Using an economic metaphor to ground intuition, x represents a sort of *poverty line*, and a “policy” S' has a better overall quality than S if some members below the poverty line in S are improved in S' , even if some of those above the poverty line in S are made worse off in S' (as long as they do not drop below the poverty line). This metaphor suggests that SE-optimality could be a reasonable criterion for specifying globally preferred solutions. We now prove that the WLO+ algorithm finds exactly the SE-optimal solutions.

Theorem 1 *The set of solutions returned by WLO+ is precisely the set of SE-optimal solutions.*

Proof:

Consider a solution S not returned by WLO+, i.e., one that is eliminated at some iteration of the WLO+ algorithm; let the optimal value (i.e., value of the weakest link) of the set of solutions be v at that iteration. Let S' be any survivor at that iteration. There must be some link i such that $U_S^i < v$ (otherwise S wouldn't be eliminated). But $U_{S'}^i \geq v$ since S' survives. Thus, $U_{S'}^i > U_S^i$. Note also that $U_{S'}^j \geq v$ for all links j . Thus, for any value k such that $U_S^k \leq v$, we have $U_{S'}^k \geq U_S^k$. It follows that S is dominated at stratum v .

Conversely, suppose S is dominated at some stratum v but, for the sake of contradiction, suppose S is not excluded from the set of solutions returned by WLO+. From the assumption that S is dominated at stratum v , there exists an S' and i such that $v > U_S^i$ and $U_{S'}^i > U_S^i$, and for any j , $U_{S'}^j \leq v$ implies $U_S^j \geq U_{S'}^j$. During the execution of the WLO+ algorithm, an increasing sequence of preference values $V = v_1, v_2, \dots, v_N = 1$ (where 1 is the “best” preference value) is created, representing the WLO optimal values at each iteration. Clearly, $U_S^i < 1$ (where 1 is the “best” preference value), so one of the V 's must exceed U_S^i . Suppose v_K is the smallest element in V such that $v_K > U_S^i$. Note that S would be removed at this iteration, as a result of its being not WLO optimal, unless the preference function for link i had been reset at an iteration $J < K$. But that function would get reset only if i was a weakest link at J . Then $v_J \leq U_S^i$ since $J < K$, and v_K is the smallest V such that $v_K > U_S^i$. Note however, that for all links j , either $U_{S'}^j \geq v > v_J$ or $U_{S'}^j \geq U_S^j$. Thus, S' would have survived to this iteration if S had. However, $U_{S'}^i > U_S^i \geq v_J$, which contradicts the fact that i is a weakest link. \square

¹This requires that the preference values be totally ordered

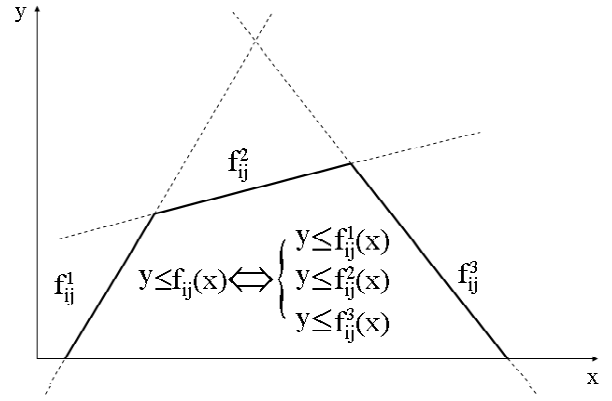


Figure 4: Convex Piecewise Linear Function

Utilitarian Optimality

Perhaps the most natural criterion for global optimality is *utilitarian*, where the global value of a solution is the sum of the local values. In this section, we consider applying a utilitarian global optimality criterion to the temporal preference problem. We show that finding a utilitarian optimal solution is tractable in the case where all the preference functions are convex and piecewise linear. Piecewise linear preference functions characterize soft constraints in many real scheduling problems; for example, in vehicle routing (where the best solutions are close to desired start times) and in dynamic rescheduling (where the goal is to find solutions that minimally perturb the original schedule).

Consider an STPP with preferences $F = \{f_{ij}\}$, and assume that the goal is to find a utilitarian optimal solution S , i.e. where $\sum_{ij} f_{ij}(S)$ is optimal. Suppose each f_{ij} is convex and piecewise linear. Thus, there is a sequence of intersecting line segments that make up f_{ij} . We will denote the individual linear functions corresponding to the segments by $f_{ij}^1, f_{ij}^2, \dots, f_{ij}^{m_{ij}}$, as illustrated in figure 4.

In this case, we will show that the utilitarian optimization problem can be reduced to a Linear Programming Problem (LPP), which is known to be solvable in polynomial time by Karmarkar's Algorithm (Cormen *et al.* 1990). This result generalizes the observation in (Khatib *et al.* 2001) that STPPs with linear preference functions can be mapped into an LPP.

Since the f 's are convex, notice that $y \leq f_{ij}(x)$ if and only if $y \leq f_{ij}^1(x) \wedge y \leq f_{ij}^2(x) \wedge \dots \wedge y \leq f_{ij}^{m_{ij}}(x)$. (See figure 4.) For the LPP, we introduce an auxiliary variable Z_{ij} , together with m_{ij} additional linear constraints of the form

$$Z_{ij} \leq f_{ij}^k(S)$$

for each f_{ij} . We also introduce a set of variables $X = \{X_1, X_2, \dots, X_n\}$ where X_i and X_j represent, respectively, the start and end points f_{ij} . An interval $[p_{ij}, q_{ij}]$ denotes the domain of f_{ij} .

The complete LPP can now be formulated as follows. The indices are assumed to range over their available values, which should be clear from the above discussion. Note

that ij in $\{f_{ij}\}$ and $\{Z_{ij}\}$ range over the edges associated with preferences. This could be a small subset of the entire edges in real applications. Finally, we introduce a variable S_{ij} for each temporal distance assignment in a solution.

- Variables: $\{X_i\}$, $\{S_{ij}\}$, and $\{Z_{ij}\}$.
- Constraints (conjunctive over all values of the indices):
 1. $S_{ij} = X_j - X_i$
 2. $p_{ij} \leq S_{ij} \leq q_{ij}$
 3. $Z_{ij} \leq f_{ij}^k(S)$
- Objective Function: $\sum_{ij} Z_{ij}$

Theorem 2 *The solution to the LPP as formulated above provides a utilitarian optimal solution to the STPP.*

Proof:

Consider the candidate STPP solution S obtained from the values of the $\{X_i\}$ variables in an optimal solution of the LPP. Clearly, the constraints (items 1 and 2) guarantee that S satisfies the STP underlying the STPP. It only remains to show that it is optimal in the utilitarian ordering for the STPP. From the constraints in item 3, we see that $Z_{ij} \leq f_{ij}^k(S)$ for each linear component k and hence $Z_{ij} \leq f_{ij}(S)$. We claim that $Z_{ij} = f_{ij}(S)$. To see this, note that the Z_{ij} variables can be varied independently without affecting the constraints in items 1 and 2. If $Z_{ij} < f_{ij}(S)$, then the objective function can be increased, without violating any constraints, by increasing Z_{ij} to $f_{ij}(S)$, which contradicts the assumption that the solution is already optimal. Thus, $Z_{ij} = f_{ij}(S)$ for each ij , and so $\sum_{ij} Z_{ij} = \sum_{ij} f_{ij}(S)$.

Suppose now there was a better solution S' for the STPP in terms of the utilitarian ordering. Then $\sum_{ij} f_{ij}(S') > \sum_{ij} f_{ij}(S) = \sum_{ij} Z_{ij}$. Observe that we can now formulate a better solution to the LPP based on S' (where we set $Z'_{ij} = f_{ij}(S')$), which is a contradiction. Thus, the S obtained from the LPP is also optimal for the STPP. \square

Example The rover STPP, shown in Figure 1, maps into the following LPP:

- Variables
 $\{cpu_1^s, cpu_1^e, ins_1^s, ins_1^e, cpu_2^s, cpu_2^e, ins_2^s, ins_2^e, T\}$,
 $\{d_{cpu_1^s, cpu_1^e}, d_{cpu_1^e, ins_1^s}, \dots\}$, and $\{z_1, z_2\}$
- Constraints

$$\begin{aligned} ins_1^s - T &= 2; & ins_2^s - T &= 9; \\ cpu_1^e - cpu_1^s &\geq 0; & ins_1^s - cpu_1^s &\geq 0; \\ cpu_1^e - ins_1^e &\geq 0; & ins_1^e - ins_1^s &= 3; \\ cpu_2^e - cpu_2^s &\geq 0; & ins_2^s - cpu_2^s &\geq 0; \\ cpu_2^e - ins_2^e &\geq 0; & ins_2^e - ins_2^s &= 1; \\ z_1 &\leq -d_{cpu_1^s, cpu_1^e}; & z_2 &\leq -d_{cpu_2^s, cpu_2^e}; \end{aligned}$$
- Objective

$$\text{maximize}(z_1 + z_2)$$

The rover example introduces only linear preference constraints, a special case of piece-wise linear functions. Notice that the z_i stand for preference values, defined as the negation of the duration assigned to the link.

Discussion and Conclusion

A potential for tradeoff emerges in the fact that WLO+ generates flexible temporal plans that are guaranteed SE-optimal but not necessarily Utilitarian optimal with respect to preferences, whereas an LPP solver of the same problem generates a fixed plan that is guaranteed to be utilitarian optimal. But if WLO+ tends to generate solutions that are always either utilitarian or “nearly” utilitarian optimal, then it may be considered a better choice for solving many problems due to its robustness in generating flexible plans. Experiments are currently underway evaluating the utilitarian quality of WLO+ solutions over a diverse set of problems, as well as to compare the run-time performance of the two formulations. The results of these experiments are too preliminary to summarize here, and will be the subject of a future report.

Another long range goal of this work is to integrate probabilities with preferences in temporal constraint reasoning, incorporating the results of (Tsamardinos *et al.* 2003). This integration will allow for a formulation of the expected utility paradigm of classical decision theory (Savage 1954) within a constraint-based setting.

The work reported here contributes to the overall goal of increasing the adeptness of automated systems for planning and scheduling. The objectives of this work overlap with those of a number of diverse research efforts. First, this work offers an alternative approach for reasoning about preferences to approaches based on multi-objective decision theory (Bacchus and Grove 1995). This work also contributes to, and builds upon, the on-going effort to extend CSP algorithms and representations to solve optimization problems or problems where knowledge is uncertain (for example, (Dubois *et al.* 1996)). Finally, the focus on solving problems involving linear piece-wise constraints has similarities to other efforts more grounded in Operations Research (for example, (Ajili *et al.* 2003)).

References

- F. Ajili and H. El Sakkout. A Probe-Based Algorithm for Piecewise Linear Optimization in Scheduling. In *Annals of Operations Research*, 118: 25–48, 2003.
- F. Bacchus and A. Grove. Graphical Models for Preference and Utility. *Uncertainty in Artificial Intelligence. Proceedings of the Eleventh Conference*, Morgan Kaufmann Publishers, San Francisco, 3–10, 1995.
- J. Bresina, K. Golden, D. Smith and R. Washington. Increased Flexibility and Robustness of Mars Rovers. *Proceedings of the 5th International Symposium on Artificial Intelligence, Robotics and Automation for Space*. ESTEC, Noordwijk, Netherlands, 1999.
- T.H. Cormen, C.E. Leiserson, and R.L. Rivest. *Introduction to Algorithms*. MIT press, Cambridge, MA, 1990.
- R. Dechter, I. Meiri, and J. Pearl. Temporal constraint networks. *Artificial Intelligence*, 49, 1991:61–95, 1991.
- D. Dubois, H. Fargier, and Henri Prade, Possibility Theory in Constraint Satisfaction Problems: Handling Pri-

ority, Preference and Uncertainty. *Applied Intelligence*, 6(4):287–309, 1996.

H. El Sakkout, M. Wallace, and T. Richards. Minimal Perturbance in Dynamic Scheduling. *Proc. of the 13th European Conference on Artificial Intelligence (ECAI-98)*: 504–508, 1998.

L. Khatib, P. Morris, R. Morris and F. Rossi. Temporal Reasoning about Preferences. In *Proceedings of IJCAI-01*. Morgan Kaufman, 2001.

L. Khatib, P. Morris, R. Morris and K. Venable Tractable Pareto Optimization of Temporal Preferences In *Proceedings of IJCAI-03*. Morgan Kaufman, 2003.

N. Muscettola, P. Morris, and I. Tsamardinos. Reformulating temporal plans for efficient execution. In *Proceedings of Sixth International Conference on Principles of Knowledge Representation and Reasoning (KR'98)*, 1998 .

L. Savage. *The Foundation of Statistics*, Dover, New York, 1954.

Paolo Toth and Daniele Vigo. *The Vehicle Routing Problem*. SIAM Monographs on Discrete Mathematics and Applications, Philadelphia, 2000.

I. Tsamardinos, M. E. Pollack, and S. Ramakrishnan. Assessing the Probability of Legal Execution of Plans with Temporal Uncertainty, In *ICAPS-04 Workshop on Planning under Uncertainty and Incomplete Information*, June, 2003.